AD-A134 036 THEORETICAL AND EXPERIMENTAL COMPARISON OF FOUR VECTOR PROCESSORS FOR NUMERICAL OCEAN MODELING(U) JAYCOR ALEXANDRIA VA A J WALLCRAFT DEC 81 1/1

UNCLASSIFIED JAYCOR-J206-81-020/6204 N00014-81-C-0085 F/G 8/10 NL

END

FILMED

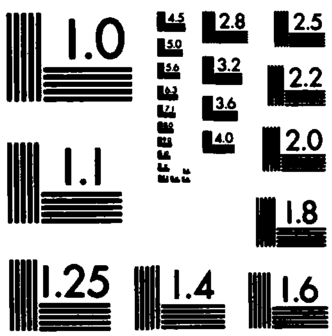

MICROCOPY RESOLUTION TEST CHART
NATIONAL BUREAU OF STANDARDS-1963-A

AD-A134036

DTIC FILE COPY

THEORETICAL AND EXPERIMENTAL COMPARISON

OF FOUR VECTOR PROCESSORS FOR

NUMERICAL OCEAN MODELING

J206-81-020/6204

JAYCOR

205 South Whiting Street
Alexandria, Virginia 22304

83 10 2

THEORETICAL AND EXPERIMENTAL COMPARISON

OF FOUR VECTOR PROCESSORS FOR

NUMERICAL OCEAN MODELING

J206-81-020/6204

Alan J. Wallcraft

Prepared For

Naval Ocean Research and Development Activity
NSTL Station, MS 39529

Under

Contract N00014-81-C-0085

December 1981

UNCLASSIFIED

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

REPORT DOCUMENTATION PAGE

READ INSTRUCTIONS BEFORE COMPLETING FORM

1. REPORT NUMBER J206-81-020/6204	2. GOVT ACCESSION NO.
3. RECIPIENT'S CATALOG NUMBER	
4. TITLE (and Subtitle) THEORETICAL AND EXPERIMENTAL COMPARISON OF FOUR VECTOR PROCESSORS FOR NUMERICAL OCEAN MODELING	5. TYPE OF REPORT & PERIOD COVERED Technical Report 11/25/80 - 11/24/81
	6. PERFORMING ORG. REPORT NUMBER J206-81-020/6204
7. AUTHOR(s) Dr. Alan Wallcraft	8. CONTRACT OR GRANT NUMBER(s) N00014-81-C-0085
9. PERFORMING ORGANIZATION NAME AND ADDRESS JAYCOR 205 South Whiting Street Alexandria, VA 22304	10. PROGRAM ELEMENT, PROJECT, TASK AREA & WORK UNIT NUMBERS
11. CONTROLLING OFFICE NAME AND ADDRESS Naval Ocean Research & Development Activity NSTL Station, MS 39529	12. REPORT DATE
	13. NUMBER OF PAGES
14. MONITORING AGENCY NAME & ADDRESS (if different from Controlling Office) Naval Ocean Research & Development Activity NSTL Station, MS 39529	15. SECURITY CLASS. (of this report) UNCLASSIFIED
	15a. DECLASSIFICATION/DOWNGRADING SCHEDULE

16. DISTRIBUTION STATEMENT (of this Report)

1 - Scientific Officer
1 - Administrative Contracting Officer
6 - NRL, Code 4627 .
12 - Defense Technical Information Center
1 - Office of Naval Research

APPROVED FOR PUBLIC RELEASE
DISTRIBUTION UNLIMITED

17. DISTRIBUTION STATEMENT (of the abstract entered in Block 20, if different from Report)

18. SUPPLEMENTARY NOTES

19. KEY WORDS (Continue on reverse side if necessary and identify by block number)

20. ABSTRACT (Continue on reverse side if necessary and identify by block number)

Super computers are required for effective ocean simulation numerical experiments because of the time and space scales of the underlying physical processes. This report compares four such computers, the Texas Instruments Advanced Scientific Computer (TIASC), the CRAY-1, the Cyber 203 and the Cyber 205 entirely on the basis of their suitability for numerical ocean modeling. Using this as the basis for comparison, it is found that the Cyber 205 is the highest rated machine.

DD FORM 1 JAN 73 1473 EDITION OF 1 NOV 65 IS OBSOLETE

UNCLASSIFIED

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

CONTENTS

1. INTRODUCTION

The Environmental Simulation Branch of the Numerical Modeling Division (Code 322) at NORDA was set up to provide a link between the numerical ocean modeling carried out by the academic community and the numerical ocean forecasting required for operational use by the Navy. To that end the branch carries out its own "academic" research and provides operational software to the Navy. In both areas state-of-the-art super computers are required for effective numerical experiments because of the time and space scales of the underlying physical processes [Hurlburt, 1981].

This report compares four such computers, the Texas Instruments Advanced Scientific Computer (TIASC), the CRAY-1, the Cyber 203 and the Cyber 205, entirely on the basis of their suitability for numerical ocean modeling. All these machines are vector processors, that is, it is only possible to attain a significant fraction of full machine speed when operating on large regularly ordered data structures, or "vectors." The exact definition of a vector varies from machine to machine, but all include one dimensional FORTRAN arrays. Therefore finite difference three dimensional ocean models (level or layer type) are vectorizable with a vector length of, at least, the number of nodes across a horizontal layer (or level). Numerical ocean modeling is an application particularly well suited to vector processors, so conclusions drawn by this report do not necessarily apply to other uses of such machines. Three of the machines are also very good scalar computers, but the TIASC has poor scalar performance and is therefore not a good general purpose machine. This fact would make the TIASC a poor choice for a university environment but has little effect on its speed in large scale oceanographic applications.

FORTRAN is the standard language for large number crunching programs, including numerical ocean modeling, and therefore all statistics (theoretical or experimental) are given for standard FORTRAN programs. On some vector processors the full power of the machine is only accessible in machine code or by using extensions to FORTRAN, this is for the most part due to the lack of sophistication of the corresponding FORTRAN compilers and so the statistics are subject to improvement as compiler software is upgraded.

2. MACHINE CHARACTERISTICS

A. Architecture

Vector operations can be divided into two phases, a start-up phase which prepares the machine for the vector operation and a solution phase which returns the results at a fixed pace per element. The start-up time is independent of vector length and can be quite long, so short vector operations take more time overall per element than operations on long vectors. A useful scale independent parameter is the vector length required to obtain a given fraction of machine speed. Taken together with the maximum vector rate (in Mflops - millions of floating point operations per second) it provides a characterization of effective machine speed.

The CRAY-1 is difficult to summarize in this way, the other machines perform vector operations memory to memory but the CRAY-1 performs such *operations vector register to vector register.* Its eight (sixty-four word) vector registers play the same role as conventional scaler registers, i.e., vector operations can be performed faster than the memory bandwidth would otherwise allow. For example, frequently used vectors can be held in the registers and temporary results need never be stored in main memory. However memory bandwidth is still the limiting factor in many situations (since all the vectors required must be transferred to registers at some time) and hence the difference between maximum possible vector speed and maximum typical vector speed (240 Mflops against 50 Mflops).

TABLE 1. MACHINE SPEEDS

Machine	No. Pipes	Word Length	Max. Typical (FORTRAN) Vector Speed	Max. Typical Vector Lengths For: 50% Speed	90% speed
TIASC	2	64	9 Mflops	40	350
		32	25 Mflops	90	800
CRAY-1	(2)	64	50 Mflops	20	50
Cyber 203	2	64	37 Mflops	150	1,400
		32	100 Mflops	400	4,000
Cyber 205	1	64	50 Mflops	50	450
		32	100 Mflops	100	900
Cyber 205	2	64	100 Mflops	100	900
		32	200 Mflops	205	1,900
Cyber 205	4	64	200 Mflops	205	1,900
		32	400 Mflops	410	3,700

Notes on Table 1.

1) The number of vector pipes is an important machine parameter, the pipes can be thought of as acting in parallel, so a 4-pipe version of a

given machine will be asyptotically twice as fast as a 2-pipe version. Differences in the number of pipes is not significant between machine types, a 2-pipe Cyber 205 is about four times as fast as a 4-pipe TIASC for example. The TIASC and the Cyber 205 can have 1, 2 or 4 pipes, the Cyber 203 always has 2 pipes, and the CRAY-1 has 12 pipes but each is dedicated to a particular operation and only the floating point addition and multiplication pipes are counted here.

2) Most numerical ocean models will perform satisfactorly with 32-bit words, which hold floating point numbers to about six significant decimal digits. The Cyber 203 and 205 have the hardware capability to process 32-bit words in vector mode, but this facility is not currently implemented in FORTRAN - it is expected that this will be added in the near future.

3) The maximum speeds quoted are those expected in a typical FORTRAN program acting on very long vectors (containing say 64,000 elements). Times for addition and multiplication are different in 64-bit mode on the TIASC and the Cyber 203, the quoted rate is for a ratio of two additions for each multiplication. The maximum speed of the CRAY-1 is highly problem dependent, ranging (even for optimized machine language codes) between 30 and 130 Mflops. The typical speed, particularly in FORTRAN, is about 50 Mflops [Temperton, 1979; Jordan and Fong, 1977].

4) All of these machines perform certain operations considerably faster than the maximum typical rate. The TIASC performs a vector dot product:

```
      P=0.0
      DO 11 I=1,L
        P=P+X(I)*Y(I)
   11 CONTINUE
```

twice as fast as conventional vector operations (e.g., 50 Mflops in 32-bit

mode), but this is not very useful in oceanographic applications.

The Cyber 205 performs an addition and a multiplication on one scalar and two vector operations such as

```
      DO 21 I=1,L
        Z(I)=X(I)+S*Y(I)
   21 CONTINUE
```

twice as fast as conventional vector operations (e.g., 800 Mflops in 32-bit mode on 4-pipe machine). This 'linked triad' capability is very useful in oceanographic applications since a significant fraction of all multiplications in ocean models are at the above form.

The CRAY-1 performs exceptionally well when a large number of vector operations are performed on a small number of distinct vectors; an equal number of additions and multiplications is also desirable. Speeds of more than 100 Mflops are obtainable in some cases, although probably not in FORTRAN. These conditions do not usually apply to ocean models.

5) Most machines achieve half speed on vectors at lengths 100 to 400 and 90% of full speed at length 1,000 to 4,000. The CRAY-1 produces a significant fraction of full machine speed on very short vectors and is therefore a better balanced machine for a general mix of programs. However, actual machine speed must also be considered, for example a 4-pipe Cyber 205 in 32-bit mode runs at the CRAY-1's maximum typical speed (50 Mflops) on vectors of length 58.

6) Since the Cyber 203 and 205 have very similar architectures it can be stated with confidence that, on any given program, the order of execution times will always be Cyber 203 (slowest), 1-pipe Cyber 205, 2-pipe Cyber 205, and 4-pipe Cyber 205. On a given Cyber 205 the 64-bit performance is identical with the 32-bit performance on the version with half as many

pipes. The 1-pipe Cyber 205 is the direct replacement for the Cyber 203; it is faster in 64-bit mode, has a lower vector start-up overhead, the linked triad capability and faster data motion operations.

Another important architectural property of these machines is their definitions of what constitutes a vector. In each case this can be characterized by a one dimensional FORTRAN array indexed with a linear combination of up to three loop index variables.

e.g.
```
      DO 31 L=LF,LL
      DO 31 J=JF,JL
      DO 31 I=IF,IL
         ...=X(K0+K1*I+K2*J+K3*L)...
   31 CONTINUE
```

where

TABLE 2.

Computer	K0	K1	K2	K3
TIASC	Integer	-1,0,+1	Integer	Integer
CRAY-1	Integer	Non-negative	0	0
Cyber 203/205	Integer	0,+1	0	0

On the TIASC the vector length is (IL-IF+1)*(JL-JF+1)*(LL-LF+1), although any of the three loops may have length 1. The definition of a vector is very general, it includes (subarrays of) three dimensional FORTRAN arrays but additionally the same element of X can appear several times in the vector, for example a matrix with constant rows could be represented as just one row.

Vector performance is degraded if the inner loop is not used, i.e., if the elements at the lowest level are not contiguous in memory. On the CRAY-1, JF=JL and LF=LL so vector length is (IL-IF+1); vectors must be accessed in ascending order, they need not be contiguous in memory but transfer to and from the vector registers may be degraded if they are not. The Cyber 203 and 205 also has vector length (IL-IF+1) but here vectors must be contiguous in memory and be accessed in ascending order. On all machines scalar variables can be treated as vectors with constant elements (i.e., K1=0 is allowed).

Each machine deals with the problem of vector overhead in a different way. On the TIASC the definition of a vector is very general so the typical length of a vector is longer on this machine than on the others, thus minimizing the effect of its quite long vector start-up time. On the CRAY-1 vector start-up time is very short, and so the definition of a vector can be less general. The Cyber 203 and the 205 have a very simple definition of a vector and a long vector start-up time, however a large selection of data motion and manipulation operations have been provided. Longer vectors can be obtained by, for example, packing non-contiguous data structures into contiguous form for vector operations and then unpacking the result, and other possibilities also exist. However many of these data motion operations are very inefficient on the Cyber 203. This machine is therefore the least flexible of those described here. On the other hand the Cyber 205 is potentially the most flexible vector machine, although this potential has not yet been realized in FORTRAN.

B. Storage

A good rule of thumb for numerical ocean models is that five to ten

grid points are required across any major features of interest (e.g., eddies, major currents, seamounts, etc.) if it is to be adequately resolved. The grid resolution required when modeling actual ocean basins can therefore be bounded by consideration of observed features. For example a grid resolution of 10 km would provide five grid points across the major seamounts in the New England chain, which have an important effect on the downstream variability of the Gulf Stream. Possible grid resolutions for several ocean regions and the corresponding storage requirements for a two-layer free surface semi-implicit hydrodynamic model, together with (very approximate) CRAY-1 computer times for a ten model year experiment, are given below in Table 3 [Hurlburt, 1981].

TABLE 3. MODEL REQUIREMENTS

Region	Grid Resolution	Grid Size	Time Step (hours)	Storage (M)	Time for 10 year run on CRAY-1 (hours)
Gulf of Mexico	10 km x 10 km	160 x 96	0.75	0.3	4
	5 km x 5 km	320 x 192	0.375	1.4	35
Western Med.	10 km x 10 km	188 x 100	0.75	0.4	6
Mediteranean	10 km x 10 km	370 x 177	0.75	1.4	20
North Atlantic	25 km x 25 km	160 x 160	1.0	0.6	5
	10 km x 10 km	400 x 400	0.5	3.5	75
World Ocean	$1^{\circ} \times 1^{\circ}$	360 x 130	1.5	1.0	8
	$0.5^{\circ} \times 0.5^{\circ}$	720 x 260	0.75	4.1	60

Actual storage requirements will vary from ocean model to ocean model, and also depend on other factors, but it is clear that realistic modeling (or forecasting) in large ocean basins, such as the North Atlantic, will require about 4 M words of storage.

Possible main memory configurations for the various machines are:

TABLE 4.

Machine	Main Memory 32-bit words	Main Memory 64-bit words
TIASC	1 M	0.5 M
CRAY-1	-	1 to 4 M
Cyber 203	2 M	1 M
Cyber 205	2 to 8 M	1 to 4 M

Both the CRAY-1 and the Cyber 205 have the potential (depending on configuration) to hold 4 M words in main memory. Even if sufficient main memory is not available it is theoretically possible to run such experiments 'out of core' by using an external storage device (usually a disk) to hold inactive arrays. The Cyber 203 and 205 have a virtual memory management system which automatically moves arrays between main storage and disk as required, however out of core ocean model calculations are not practical on these machines for reasons detailed below in the discussion of ocean forecasting. On the CRAY-1 and TIASC the movement to and from disk must be

implicitly controlled by the program, in the best case disk I/O is performed entirely in parallel with computations and the code runs as if it were core contained. But even if this best case, which may not be attainable in practice, the computing time required to execute these large models on the CRAY-1 (or the slower TIASC) is prohibitive. If it is assumed that the practical limit on computing time is about ten hours for a ten year model run then an approximate upper limit on model storage requirement can be determined.

TABLE 5.

Machine	Pipes	Max. Storage per model 32-bit	64-bit
TIASC	2	0.6 M	0.3 M
CRAY-1	(2)	-	1.0 M
Cyber 203	2	2.0 M	1.0 M
Cyber 205	2	3.0 M	2.0 M
Cyber 205	4	5.0 M	3.0 M

Table 5 does not necessarily indicate the optimal main memory configurations for several reasons:

1) Different models have different storage and computer time requirements; however the example model is of an efficient design.

2) Ten hours of computer time may be an overestimate of the time available

for an experiment.

3) Out of core calculations are possible in the TIASC and CRAY-1.

4) The model will probably run in a timesharing environment, so the full machine may not be available.

5) Storage can be traded off against execution time, in particular the most efficient methods for solving a Helmholtz's equation require more storage than has been allowed here.

However it is clear that only the Cyber 205 is potentially fast enough for realistic long time scale modeling of large ocean basins.

The requirements of ocean forecasting are a little different. The length of a forecast is measured in days (or months) rather than years and the model will probably run in stand alone mode so the full machine will be available, but it is real time, rather than computer time, which is the important parameter here. In the development stage several long time scale experiments will be required to test the model, which will also have to be spun up before the first forecast. The CRAY-1 and TIASC are almost certainly too slow to allow the development of such a forecasting model with satisfactory grid resolution.

The Cyber 203 and 205 have a virtual memory system and it might be supposed that, since the forecast is over a short time scale, the model could run out of core. As a counter example consider a model requiring 4 M words of storage executing on a machine with 2 M words of memory. Since all values are accessed every time-step an absolute minimum of 2 M words must be swapped into main memory per time-step. Variables are moved into memory in units of pages, and 2 M words take up 32 large pages, so at least 32 page faults will occur per time-step. The process of swapping in a new large page takes about

half a second of wall clock time (and a very small amount of computer time) so the hypothetical model would spend a minimum of about 16 seconds each time-step in page faults. This figure would not be achieved in practice, 60 seconds of page fault time per time-step would be more realistic and at this value the model would take about one hour for a 3 day forecast (assuming a time-step of one hour). The same forecast running in core on a 4-pipe Cyber 205 might take 20 seconds. Similar arguments demonstrate that long time scale ocean models must also be memory resident (e.g., a 4 M word experiment taking 10 hours of computer time might have a turn around time of one month on a 2 M word machine).

It is clear that a forecast model requiring the maximum configuration of 8 M (32-bit) words is practical on the 4-pipe Cyber 205. However there is little existing experience in ocean forecasting with high horizontal resolution and it is not clear that such a model would be useful given the state-of-the-art in real time ocean data collection and assimilation. The quantity and quality of data available is expected to increase rapidly, particularly satellite data, and therefore by the mid 80's a need might well exist for a forecasting model of such a size. Of course, by then machines even faster than the Cyber 205 might be commercially available. NORDA is currently developing a World Ocean Model to run on the Cyber 203 (and therefore in 2 M 32-bit words). Treating the world ocean as three separate oceans might be one possibility (at least in this case) for maximizing grid resolution in a given amount of memory.

C. Software

FORTRAN is not a good vector programming language; arrays are second class objects that can only (usually) be accessed element by element, often

within 'DO' loops. NORDA's approach to using vector processors is to write standard FORTRAN programs in such a way that a 'vectorizing' FORTRAN compiler can recognize the underlying vector structure of such 'DO' loops and produce vector code where appropriate. The alternative approach, of using non-standard extensions to FORTRAN or even coding in assembly language, is not acceptable at NORDA because its products must be transportable. Standard FORTRAN programs are also easier to understand and to modify, important properties for ocean models, since minor changes to the code are made routinely when developing a version of the model suitable for a given ocean region.

Some manufacturers strongly advocate the use of vector extensions to FORTRAN, arguing that it is not possible to vectorize all FORTRAN codes [Kasic, 1979; Mossberg, 1981]. It is certainly true that a FORTRAN code written for a scalar machine may be inefficient on a vector processor. But if a code is written from scratch for a vector machine in, possibly highly stylized, standard FORTRAN then the full power of the vector architecture should be available via a good vectorizing compiler. The vector extension approach has two advantages for the manufacturer: it provides a strong incentive to remain within a computer family when upgrading a system and it relieves the pressure to commit resouces to the development of a good vectorizing compiler. On the other hand it is not obvious that a code written in FORTRAN to vectorize on one machine will necessarily vectorize on a different vector processor. However an ocean model written in FORTRAN to vectorize on the TIASC was transferred to the Cyber 203 in one man-day, and a fully vectorizing version was produced within one man-week [Wallcraft, 1981]. If the Cyber 203 had a good vectorizing compiler the transfer would

have been completed in one man-day, but if the original version had been written using TIASC vector extensions to FORTRAN, then producing a version using Cyber 203 vector extensions might have taken several man-months.

The quality of existing vectorizing FORTRAN compilers differs from machine to machine:

1) TIASC

The most sophisticated compiler currently available. It will vectorize almost all theoretically vectorizable nests of up to three loops. It is not usually possible to produce any significant improvement in speed by using vector extensions to FORTRAN or assembly language.

2) CRAY-1

A good inner loop vectorizer, which is sufficient given the machines efficiency on short loops. In some cases a significant improvement in speed is possible by using CRAY assembly language.

3) Cyber 203

A poor inner loop vectorizer is coupled with a very limited ability to vectorize outer loops. None of the machines extensive collection of manipulation operations are available (either implicitly or explicitly) via standard FORTRAN. In many cases a very significant improvement in speed is possible using vector extensions to FORTRAN.

4) Cyber 205

Similar to the Cyber 203 except that inner loops with non-unit incrementation parameters are vectorized, and linked triad operations recognized.

The vectorizing compilers on the Cyber 203 and 205 are less well developed than those on the other two machines. Their inner loop

vectorizer is significantly less spohisticated than that available on the CRAY-1, and in any case inner loop vectorization is not sufficient given the long vector start-up times of those machines. The Cyber 205 has a very efficient implementation of a very flexible vector architecture. For example, the TIASC vector architecture would be very efficiently emulated on the 205. This means that techniques introduced at least six years ago for outer loops vectorization on the TIASC [Wedel, 1975] are equally applicable to the Cyber 205, and there is therefore no excuse for the poor performance of the Cyber 205 compiler. The Cyber 205 is a new machine and it is likely that the vectorizor will be substantially improved in the future. Relatively minor improvements in some areas would have a large effect on the machine's FORTRAN performance on ocean models. The Cyber 203 has been superseded by the 205 and improvements to this machine's FORTRAN performance are less likely, particularly since many of its data motion operations are very slow.

Another major deficiency of the FORTRAN compiler on the Cyber 203 and 205 is that REAL variables are stored in 64-bit words. This size was probably chosen for compatability with other CDC machines, but it effectively reduces the speed of the vector processor by half (or more on the 203) since 32-bit arithmetic is not available in any practical way to the FORTRAN programmer, not even by using FORTRAN vector extensions. A compiler with 32-bit capability has been promised by CDC but its exact form is not known. The best solution (for oceanographers) would be to redefine REAL variables as 32-bit words, 64-bit DOUBLE PRECISION variables would then also be vectorizable. An acceptable alternative would be to introduce a new type, say REAL*4, and allow it to be used interchangeably with other types. Automatic vectorization must apply to the new type and an IMPLICIT statement would be

useful. A minimal solution, which is absolutely not acceptable, might be to introduce the REAL*4 type but only allow its use within vector extensions to FORTRAN.

Other areas of system software will not be considered here since the CRAY-1, Cyber 203 and 205 require a front end processor which will provide the major user interface to the operating system. (The TIASC has an IBM based operating system.) Applications packages, for linear algebra or statistics for example, are also important but are usually provided by users of the machines. The CRAY-1 has a good range of such software as does the TIASC although its quality is somewhat variable on the latter machine. The Cyber 203 and 205 have packages originally written for the STAR computer. The Cyber 205 now has a large user base and application software specifically for this machine can be expected in the near future.

3. EXPERIMENTAL COMPARISONS

A. A Reduced Gravity Ocean Model

Model execution times are presented for a one layer reduced gravity ocean model set up for experiments on a rectangle representing the Gulf of Mexico [Hurlburt and Thompson, 1980]. The model is free surface, primitive equation, treats gravity waves implicitly, neglects thermodynamics, and is written entirely in standard FORTRAN. The execution time per model year is given for two mesh sizes, 80 x 48 and 160 x 96, with timesteps of 90 minutes and 45 minutes respectively (these timesteps are not maximal, they were used in the Gulf of Mexico experiments for compatibility with results from other models). The execution times are subdivided into two parts, the time expended in calculating the solution to the Helmholtz's explicit equation required each timestep (the solver time) and everything else (the explicit time). This subdivision together with the fact that the explicit time is for 65 additions, 36 multiplications and 2 divisions (with 22 linked triads) at each mesh node allows similar tables to be drawn up for other ocean models based on the data presented here.

Times on the TIASC and the Cyber 203 in 64-bit mode were obtained from actual computer runs. Times for the CRAY-1 were estimated from published solver times [Temperton, 1979] and from computer runs of a two layer quasi-geostrophic model [Chow, 1981]. Times for the Cyber 203 in 32-bit mode and for the Cyber 205 were estimated from a detailed breakdown of the 64-bit Cyber 203 times. These estimates are thought to be very accurate (say within 5%) because each machine has the same scalar processor and vector times are deterministic, i.e., given the times for vector operations of known length on one machine times for a similar machine with different vector speeds can be

calculated reliably. Times for the Cyber 203 scalar box are said to be about one and a half times as fast as that on a CDC 7600, the state-of-the-art in scalar processors (represented by the AMD 470/V12) is about twice this speed but the Cyber 203 still has one of the fastest scalar processors available.

TABLE 6.

Times For a One Layer Reduced Gravity

Semi-Implicit Ocean Model on an 80 x 48 Rectangular Ocean

Computer	No Pipes	Word Length (bits)	Time Per Model Year (sec) Solver	Explicit	Total	Time Ratios S	E	T
Cyber 203/205	Scalar	64	172	360	532	15.6	45.0	28.0
TIASC	2	32	54	113	167	4.9	14.1	8.8
Cyber 203	2	64	55	74	129	5.0	9.2	6.8
CRAY-1	(2)	64	23	66	89	2.1	8.3	4.7
Cyber 203	2	32	42	32	74	3.8	4.0	3.9
Cyber 205	2	64	16	23	39	1.5	2.9	2.1
Cyber 205	2	32	13	13	26	1.2	1.6	1.4
Cyber 205	4	64	13	13	26	1.2	1.6	1.4
Cyber 205	4	32	11	8	19	1.0	1.0	1.0

TABLE 7

Times for a One Layer Reduced Gravity

Semi-Implicit Ocean Model on an 160 x 96 Rectangular Ocean

Computer	No.Pipes	Word Length (bits)	Time Per Model Year (secs) Solver	Explicit	Total	Time Ratios S	E	T
Cyber 203/205	Scalar	64	1514	2864	4378	29.7	54.0	42.1
TIASC	2	32	369	886	1255	7.2	16.7	12.1
Cyber 203	2	64	290	560	850	5.7	10.6	8.2
CRAY-1	(2)	64	165	530	695	3.2	10.0	6.7
Cyber 203	2	32	196	223	419	3.8	4.2	4.0
Cyber 205	2	64	92	173	265	1.8	3.3	2.6
Cyber 205	2	32	67	93	160	1.3	1.8	1.5
Cyber 205	4	64	67	93	160	1.3	1.8	1.5
Cyber 205	4	32	51	53	104	1.0	1.0	1.0

The Helmholtz solver used is an implementation of FACR(0) [Hockney, 1970] written in standard FORTRAN for vector machines. This algorithm is certainly the fastest known for this problem on the TIASC and the CRAY-1, it is probably also the fastest on the Cyber 203 and 205; on scalar processors FACR (1) with an optimal choice of 1 would be slightly faster. The average inner loop vector length is equal to the first dimension of the mesh (i.e., 80 or 160) and this is the actual vector length on all the machines except the TIASC which also vectorizes the outer loop and has an average vector length

about four times as long as the other machines (the outer loop typically passes over only a small number of non-contiguous values). Relative to maximum machine speed the CRAY-1 is the most efficient, with the TIASC a close second. However the Cyber 205 (with 2 or 4 pipes) is always actually faster than the CRAY-1, its basic maximum speed advantage outweighting the relative efficiency of the CRAY-1. The Cyber 203 has a very long vector start-up time (hence the difference between the times of the 203 in 32-bit mode and the 2-pipe 205 in 64-bit mode) and vectors times comparable to the, theoretically slower, TIASC on the smaller problems. The Cyber machines perform significantly better on the larger problem, both in actual speed and relative to the TIASC and CRAY-1. Solver times might be reduced 30-40% on the CRAY-1 by using an assembly language code. Times on the Cyber 205 might be reduced by rewriting the FORTRAN version to take full advantage of linked triads, but most of the time is currently spent in the vector start up phase and the present code would run significantly faster (particularly on the 4-pipe machine) if the FORTRAN compiler performed outer loop vectorization.

The vector length for the explicit section of the code is approximately the mesh dimension (3,840 or 15,360), except on the CRAY-1 which only vectorizes inner loops (length 80 or 160). Outer loop vectorization, in FORTRAN, is only possible on the Cyber 203 and 205 at the expense of additional scalar code [Wallcraft, 1981] accounting for 3 seconds on the smaller and 12 seconds on the larger problem. With such long vectors the times closely reflect each machine's maximum speed. The model contains a large number of linked triad operations which add to the Cyber 205 speed, and this is the cause of the difference between the times on the Cyber 203 in 32-bit mode and the 2-pipe 205 in 64-bit mode. If the Cyber 203 and 205 FORTRAN

compilers were improved to allow outer loop vectorization without the addition of extra scalar code the time ratios would be, 4-pipe 205 in 32-bit mode: Cyber 203 in 32-bit mode: CRAY-1: TIASC: Cyber 203 in scalar mode - 1 : 5 : 13 : 22 : 70, and the Cyber 205 speed would be over 450 Mflops.

The total execution time on the TIASC is about twice as long as on the CRAY-1, which has times between those for 64-bit and 32-bit models on the Cyber 203. The Cyber 205 is between two and seven times as fast as the CRAY-1 depending on the problem size, machine and precision under consideration. The 4-pipe Cyber 205 in 32-bit mode is at least 50 times as fast on this model as most scalar machines, it is probably 15-20 times as fast as an AMD 470/V12.

In terms of operation counts the solver phase should account for about 30% of the total execution time, but on the Cyber 203 and 205 this phase is more significant and can account for up to 60% of the total time. The relative performance of all the machines on other ocean models will therefore depend on the percentage of times expected to be used in solving elliptic partial differential equations. Fully explicit models have no solver phase and will be very efficient on the Cyber 205, as will some level type models which only require one stream-function determination per timestep. On the other hand the addition of the capability to use non-rectangular ocean basins would at least double the time spent in the solver phase. However the Cyber 205 will always be faster than the CRAY-1 (and the TIASC), even on medium sized problems (e.g., 80 x 48 mesh) and becomes relatively more efficient on the very large problems for which the machine was designed.

B. Saturation Vapor Pressure Calculation

Ocean models which include thermodynamic effects give rise to

calculations which are only conditionally performed. Because the conditionality destroys the very regular structure associated with vectors such calculations are one of the classical examples of 'non-vectorizable' code. The saturation vapor pressure calculation, taken from an atmospheric forecast model at FNOC, is of this type since one of two possible sixth order polynomials of the temperature is returned at each node depending on the temperature regime.

On a scalar computer the code might be:

```
      SUBROUTINE SATUPR
      PARAMETER (L=10000)
      COMMON/SUP/ QS(L),T(L),A0,A1,A2,A3,A4,A5,A6,
     +                       B0,B1,B2,B3,B4,B5,B6
C
      DO 11 I=1,L
        TI=T(I)
        IF(TI.LE.224.) QS(I)=A0+TI*(A1+TI*(A2+TI*(
     +                       A3+TI*(A4+TI*(A5+TI*A6)))))
        IF(TI.GT.224.) QS(I)=B0+TI*(B1+TI*(B2+TI*(
     +                       B3+TI*(B4+TI*(B5+TI*B6)))))
   11 CONTINUE
      RETURN
      END
```

On a vector processor both calculations are performed on each element and the required solution is then chosen:

```
      DO 11 I=1,L
        QS(I)=A0+T(I)*(A1+T(I)*(...))
        QT(I)=B0+T(I)*(B1+T(I)*(...))
   11 CONTINUE
      DO 12 I=1,L
        IF(T(I).GT.224.) QS(I)=QT(I)
   12 CONTINUE
      RETURN
      END
```

The vector version does twice as much work as the original but runs at vector speed. Loop 12 will not automatically vectorize on most machines so non-standard code must be used, this is of little importance here since separate scalar and vector versions must be maintained for full

transportability in any case. The Cyber 203 and 205 vector instruction set is sufficiently rich to allow the 'scalar' version to vectorize directly. However this is far beyond the capabilities of the existing FORTRAN compiler.

The routine was originally chosen for its fast execution time on the CRAY-1 [Wellck, 1981] and the original CRAY-1 times are used here. Times on the TIASC and Cyber 203 in 64-bit mode are also for actual computer runs, all other times are estimated as in the previous section.

TABLE 8

Calculation of the Saturation Vapor Pressure

Method - 6th order polynomial approximation of QS(T)

Depending on Temerature Regime (i.e. $T > 224.0$)

Vector Length	2 pipe TIASC 32-bit	CRAY-1 64-bit	203 64-bit	203 32-bit	2 pipe 205 64-bit	2 pipe 205 32-bit	4 pipe 205 32 bit
	Time Per Result (μs)						
Scalar	9.56	2.24	2.42	2.42	2.42	2.42	2.42
10	10.82	0.78	3.53	3.53	3.53	3.53	3.53
20	7.31	0.56	3.25	3.25	2.23	2.15	2.11
50	3.86	0.35	2.67	2.67	0.99	0.91	0.87
100	2.60	0.31	2.16	1.67	0.57	0.49	0.45
200	1.98	0.29	1.45	0.97	0.37	0.29	0.25
500	1.60	0.27	1.04	0.54	0.24	0.16	0.12
1,000	1.48	0.27	0.90	0.40	0.20	0.12	0.08
2,000	1.42	0.27	0.84	0.34	0.18	0.10	0.06
5,000	1.38	0.26	0.79	0.29	0.17	0.09	0.05
10,000	1.37	0.26	0.77	0.27	0.16	0.08	0.04

← estimated →

Notes on Table 8:

1) Two-pipe Cyber 205 times in 32-bit mode are identical to 4-pipe 205 times in 64-bit mode (not shown).

2) Quoted scalar times are for vector length 10,000 (i.e., subroutine call overhead is not included). Some other table entries are also scalar

times, in these cases vector times are longer.

3) A large amount of arithmetic is performed on a small amount of data. Much of the calculation executes at register to register speed in scalar mode on all machines and also on the CRAY-1 in vector mode.

4) Almost all the arithmetic can be performed as linked triads on the Cyber 205.

5) Most of the variation in CRAY-1 times with vector length is due to subroutine call overhead.

Only the results on very long vectors (5,000 or 10,000) are relevant to ocean modeling applications. The CRAY-1 is executing more than twice as fast as it does on more typical codes but it is still not significantly faster than the Cyber 203 in 32-bit mode (on long vectors). It is, however, five times as fast as the TIASC and three times as fast as the 203 in 64-bit mode. The Cyber 205 is always faster than the CRAY-1 on long vectors, the 4-pipe version in 32-bit mode is six times faster. If operations actually performed are counted the CRAY-1 is executing at about 100 Mflops and the Cyber 205 at 600 Mflops. These rates reduce to 50 and 300 Mflops if only the required operations are counted, compared to about 1 M flop on the TIASC and about 5 Mflops on the CRAY-1, Cyber 203 and Cyber 205 in scalar mode.

For non-oceanographic applications, with short vector lengths, the CRAY-1 becomes relatively more efficient. This example is not typical but it is clear that the CRAY-1 can achieve a significant fraction of machine speed on very short vectors. Scalar speeds are comparable on the CRAY-1 and Cyber 205, the 205 is much faster on long vectors, but there is a range of vector lengths over which the CRAY-1 is superior. On this example the range is about 2 to 200 for the Cyber 205 (and 2 to 10,000 for the Cyber 203). These are

probably very nearly best case figures for the CRAY-1, more typical values are given by the vector length at which the machine achieves a speed of 50 Mflops (the typical maximum CRAY-1 speed). Vector lengths for 50 Mflops are:

TIASC	= maximum speed 25 Mflops
203, 2-pipe, 64-bit	= maximum speed 37 Mflops
203, 2-pipe, 32-bit	= 400
205, 2-pipe, 64-bit	= 100
205, 2-pipe, 32-bit	= 68
205, 4-pipe, 64-bit	= 68
205, 4-pipe, 32 bit	= 58

The CRAY-1 is therefore typically faster than the Cyber 205 on vectors of length 2 to 70.

4. CONCLUSIONS

The Cyber 205 is by far the best computer currently available for numerical ocean modeling. It is the only machine with the capability to run long time scale high horizontal resolution numerical experiments on the models of realistic ocean basins which will become increasingly important in the 1980's. Two possible configurations are a 2-pipe version with 2 M (64-bit) words of storage or a 4-pipe machine also with 2 M (or possibly 3 M) words. A 2-pipe machine with 1 M words might also be just viable, but for ocean forecasting applications the 4-pipe machine is the best choice, either with 4 M words or with 2 M words and the option to upgrade to 4 M words at a later date. The FORTRAN compiler on the Cyber 205 is not at an acceptable standard and an undertaking should be sought, by any potential purchaser, from CDC on specific improvements (with delivery dates) in this area. Two improvements of particular importance to numerical ocean modeling applications are:

1. The ability to access 32-bit words in FORTRAN.
2. A full outer loop vectorization capability.

Further details are to be found in the section on software.

The CRAY-1 and Cyber 203 are of approximately equal capability for oceanographic problems. Both can be used to perform acceptable numerical experiments, but the very new Cyber 205 can be four to ten times faster on typical ocean models and this machine will therefore be used to produce the state-of-the-art numerical ocean experiments in the next few years. The TIASC is the oldest machine type considered here and the slowest. However a 4-pipe version would be comparable, in 32-bit mode, with the CRAY-1 on ocean problems, and the 2-pipe TIASC available to NORDA has allowed the Numerical Modeling Division to remain competitive through the late 1970's.

The CRAY-1 may still be the fastest machine on a general mix of programs, as might be found in a university environment. It is particularly fast at compiling FORTRAN programs for example. The rationale behind obtaining a vector processor must, at least in part, be such a machine's performance in large problems and here the Cyber 205 is outstanding. If either vector processor is front ended by a good scalar machine, such as the Cyber 175, then small jobs can execute efficiently on this machine (with fast turn around time in timesharing mode) since large jobs will be queued to the vector processor. Therefore even in a university computer environment the Cyber 205 may be the best overall choice.

In the field of super computers the most recently introduced machine is usually the fastest and there is always the temptation to wait for the next, even faster, machine to become available. However, any new machine would have to run at about 1500 Mflops (in 32-bit mode) on ocean models to offer a significant improvement in performance in this area over the 4-pipe Cyber 205. Alternatively a new machine might be comparable in speed on large problems with the 205 but be a better choice overall because of its improved performance on short vectors. But there are dangers inherent in new computer design, many proposed super computers never reach the production stage and the software support for new machines is often very poor.

To conclude, the supercomputer market place is now in the healthy state of having competing products. The Cyber 205 is the fastest machine available but the CRAY-1 can still be the best choice for some applications, although this is at least partially due to the inadequate software support available for the 205.

5. REFERENCES

Chow, J., 1981. Private Communication

Hockney, R. W., 1970. "The Potential Calculation and Some Applications", Meth. Comp. Phys. 9 pp. 135-211.

Hurlburt, H. E., 1981. "Computing Requirements for Navy Ocean Modeling". NORDA position paper.

________, and J. Dana Thompson, 1980. "A Numerical Study of Loop Current Intrusions and Eddy Shedding," J. Phys. Oceanogr. 10, pp. 1611-1651.

Jordan, T. L., and K. Fong, 1977. "Some Linear Algebraic Algorithm and Their Performance on CRAY-1" in [Kuck, Lawrie and Sameh, 1977].

Infotech Ltd., 1979. "Supercomputers" State-of-the-Art report.

Kasic, M. J., 1979. "Vector Processing on the Cyber 205" in [Infotech, 1979].

Kuck, D. J. D. H. Lawrie, and A. H. Sameth, (eds), 1979. "High Speed Computer and Algorithm Organization", Academic Press.

Mossberg, B., 1981. "An Informal Approach to Number Crunching on the Cyber 203/205," Control Data Corporation.

Temperton, C., 1979. "Fast Fouirier Transforms and Poisson Solvers on CRAY-1" in [Infotech, 1979].

Wallcraft, A. J., 1981. "Transferring Ocean Models From the TIASC to the Cyber 203" JAYCOR Report J206-81-016/6204.

Wedel, D., 1975. "FORTRAN for the Texas Instruments ASC System" SIGPLAN 10, pp. 119-139.

Wellck, R. E., 1981. "An Atmospheric Regional Forecast Model Benchmark for the CRAY-1", CRAY Research presentation.

APPENDUM

Details of the CRAY-2 have recently been announced (Datamation, Jan 1982). It will consist of four processors running in parallel each with three times the power of the CRAY-1 for a total vector speed about twelve times that of the CRAY-1. Scalar speed will be about six times the CRAY-1 and the maximum main memory capacity will be 32 M words. The machine is to be 'phased in' over the next three years, but within this time scale it is not clear when the first machine will be delivered. The CRAY-1 will continue in production for the foreseeable future; an upgraded version, the CRAY-1X, is under development and, judging by the CRAY-2 performance, this may be two or three times as fast as the current CRAY-1S.

The CRAY-2 will be capable of rates in the 1000 to 1500 mflop range for many applications and, with its corresponding improved scalar speed, will certainly be the fastest general purpose scientific number cruncher, perhaps four to ten times as fast as the 4-pipe Cyber 205.

In numerical ocean modeling applications the speed of the CRAY-1 is limited by the register to memory bandwidth. This bottleneck might be more or less severe on the CRAY-2 and so it is impossible to make totally reliable comparisons without benchmark data. However, using the figure of twelve times a CRAY-1 and the data in Tables 6 and 7, it is estimated that the CRAY-2 (in 64-bit mode) is only about as fast as a 4-pipe Cyber 205 in 32-bit mode on explicit model code, but is three to five times faster at solving elliptic PDEs. Overall the CRAY-2 might be about twice as fast as the Cyber 205 on large scale ocean models. But this figure could be in error by a factor of two either way because of the uncertainties in CRAY-2 performance and because solver times on the Cyber 205 are subject to improvement.

The CRAY-2 may not be available for several years and delivery dates are notoriously optimistic in the computer industry. However, the machine is sufficiently advanced that a potential supercomputer customer might well be tempted

to wait, particularly if an upgrade from the CRAY-1 is being considered. On the other hand, the available information indicates that an ocean modeling group with access to a 4-pipe Cyber 205 can remain competitive throughout the 1980s. Groups limited to super computers developed in the 1970s (the TIASC, CRAY-1S and Cyber 203) will be at a disadvantage by 1985 and intermediate machines (the CRAY-1X and 2-pipe Cyber 205) can only be considered stop-gap machines if state of the art ocean modeling is the goal.

CDC may soon after a second level of solid state memory, between central memory and disk storage, for the Cyber 205 (Levine, R. D. - Supercomputers - Scientific American, Jan. 1982). If this allows the virtual memory system to operate effectively on large time dependent problems, then it will considerably increase the cost effectiveness of this machine. Viable configurations for ocean modeling or forecasting might include a 2-pipe Cyber 205 with 1M (64-bit) words of central memory and 2 to 4 M words of second level memory and a 4-pipe Cyber 205 with 2 M words of central memory and 4 to 8 M words of second level memory.

END

FILMED

11-83

DTIC